Responding to Dynamic Workloads and Varying Harvested Energy in Energy Constrained Systems

Yousef Shakhsheer
Department of Electrical and Computer Engineering
University of Virginia

A Dissertation Proposal Presented in Partial Fulfillment of the
Requirement for the
Doctor of Philosophy Degree
in Electrical Engineering

July 18, 2012

Abstract

Energy efficiency is the most critical metric in many modern integrated circuits. The motivations to lower energy spans across a variety of application areas and performance requirements. Examples of this include enabling battery-less operation in energy harvesting body sensor networks (BSNs), maximizing battery life in mobile platforms, and managing processor as they heat up during period of heavy use.

Circuit designers optimize their circuits at design time and/or run time to improve energy efficiency. Design time optimization is a well-covered topic. Runtime optimizations have not been as well researched. There are many conditions that change dynamically at runtime in low performance and medium performance electronics that effect energy efficiency. These changes include the amount of data being sampled/inputted and environmental changes. These dynamics provide opportunities to save energy. Design time optimizations are unable to adapt to these runtime changes. Instead, design time optimizations have to guard band for real time changes. The inability to adapt to changes in these conditions can result in wasted energy, shorter electronics lifetimes, and device death in energy constrained systems. This thesis proposes focusing on two important system conditions that change dynamically at runtime: processing workload requirements and the amount of energy harvested in energy harvesting systems.

To allow electronics to adapt to varying workload rates, designers utilize voltage as a knob in a scheme called dynamic voltage scaling (DVS). As the workload decreases, the circuit voltage is decreased resulting in a linear degradation in speed but a quadratic energy savings. A method called Panoptic ("all-inclusive") Dynamic Voltage Scaling (PDVS) extends DVS to a finer granularity in space and time, allowing for much more flexible and energy efficient design. This work will apply PDVS to a DSP data-flow processor to show energy savings over multiple benchmark workloads as well characterizing the overheads of PDVS.

To enable a BSN node capable of operating solely from a small amount of harvested energy, BSN nodes must consume ultra-low power (10s of μWs). The BSN architecture is instrumental in achieving low power consumption. BSN architecture must allow the node to acquire, process, and transmit biosignal data while consuming low power. This work will investigate BSN architecture decisions, such as tradeoffs between

custom controls and generic microcontrollers, to enable low power operations to run solely off of harvested energy and the ability to support multiple biosignals.

To adapt to varying amounts of harvested energy, BSN nodes require a power management system specific to energy harvesting needs. This power manager must adjust node power consumption to the changing power harvesting input to ensure long term functionality. This work will demonstrate for the first time an implemented, closed loop power management system capable of adjusting node power consumption to the amount of energy harvested as well as explore the power management design space.

1. Introduction

Energy efficiency is the most critical metric in many modern integrated circuits. The motivations to lower energy ranges over a range of electronic performance such as low, medium, and high performance. This motivation includes enabling battery-less operation in energy harvesting body sensor networks (BSNs) for potentially an indefinite electronic lifetime, maximizing battery life in mobile platforms such as cell phones, and managing thermals in processor cores (Figure 1). While energy efficiency is important in all applications, this thesis focuses on low performance and medium performance energy constrained electronics due to their lifetime constraints.

To improve energy efficiency, circuit designers optimize their circuits at design time (fabrication) and/or runtime (post-fabrication) to improve energy efficiency. Design time optimization is a well-covered topic. In designs that feature only design time optimizations, designers guard band their design at fabrication time to adapt to runtime changes, but this proves to be inefficient in most cases as they have to account for the worst case to ensure robust operation. Runtime optimizations provide the possibility to adapt to these changes over time and maximize energy efficiency post fabrication.

There are many conditions that change in low performance and medium performance electronics that are capable of effecting energy efficiency. The changes could be due to the amount of data being sampled/inputted, environmental changes, or many other factors. These dynamics provide opportunities to save energy for electronics that are capable of runtime adaption. Inability to adapt to changes in these conditions can result in wasted energy, shorter electronics lifetimes, and electronics death in energy constrained systems. This thesis focuses on two important system conditions that change at runtime: processing workload requirements and the amount of energy harvested in energy harvesting systems. Adapting to runtime changes in workload and harvested energy will improve the system's overall energy efficiency.

Figure 1. Low, medium, and high performance electronics all require energy efficiency to extend lifetime or reduce heat dissipation.

Many systems occasionally require high performance, but their varying workload requirements remain below this maximum workload requirement for the majority of their lifetimes. Designing the system in a static fashion to support this peak workload can substantially increase the total system power and greatly shortens electronics lifetime. This provides an opportunity for energy savings.

First, to allow electronics to adapt to varying workload rates, designers are able to adjust voltage to change delay/performance and reduce energy in a scheme called dynamic voltage scaling (DVS). As workload decreases, the circuit voltage is decreased resulting in a linear degradation in speed but a quadratic energy savings. A method called Panoptic ("all-inclusive") Dynamic Voltage Scaling (PDVS) extends DVS to finer granularity in space and time, allowing for much more flexible and energy efficient design. PDVS allows the system to adapt to varying normalized workloads from ~0 to 1. This concept has been proven on a test chip that applied PDVS to an adder and multiplier, but has not been implemented on the system level. This concept will be applied to a DSP data-flow processor in this work to show energy savings over multiple benchmarks as the workload as well characterizing the overheads of PDVS.

Second, harvesting ambient energy, such as from the sun, vibrations, or thermal gradients due to body heat, is an appealing alternative to battery based operation, allowing for an indefinite electronics lifetime. However, energy harvesters' power output is normally low (10's of µWs [1]) and highly environment-dependent. For example, solar harvesting provides reduced power in cloudy conditions. Designing the system in a static fashion to assume a constant amount of energy harvested can result in node death or limited functionality.

To enable a BSN node capable of operating solely off of harvested energy, BSN nodes must consume low power (<~40µW). BSN architecture is instrumental in achieving low power consumption. BSN architecture must allow for flexibility to acquire, process, and transmit biosignal data while consuming little power. To be capable of consuming this little amount energy, the node should reduce the number of bits transmitted to allow the power hungry radio to be duty cycled through processing data on-node. This work will show a fabricated 130nm wireless BSN node system-on-chip (SoC) that provides real-time ECG (electrocardiogram), EEG (electroencephalography), and EMG (electromyography) acquisition, signal analysis and processing, and wireless communication over a radio. The node will be battery-less, powered entirely from a wearable thermoelectric generator (TEG) with an off-chip storage cap and an on-chip boost converter. This work will also investigate BSN architecture decisions to enable low power operations such as exploring the tradeoffs between using a generic microcontroller (MCU) and a custom controller, bus architecture between circuit switched and address switched architectures, and the system level effects of optimizing NOP operations.

To adapt to varying amounts of harvested energy, BSN nodes require a power management system specific to energy harvesting needs. This power manager must adjust node power consumption to the changing power harvesting input to ensure functionality. The power manager must be capable of managing high power operations, such as transmitting data wirelessly, which can consume 100s of µWs and may exceed the power budget set by power harvesting. This work will show an implemented, closed loop power management system capable of adjusting node power consumption to the amount of energy harvested and explores several key power management design decisions.

The contributions of this dissertation will be:

- First implementation of PDVS on a processor
- An evaluation of system level savings and overheads of PDVS for various rate vs. time profiles
- The first complete of the art BSN node SoC capable of running solely off harvested energy

- A BSN system level exploration of tradeoffs between using generic MCU and custom controller, bus architecture, and optimization of NOPs
- An energy harvesting-specific power management system capable of responding to changes in the amount of energy harvested
- An exploration of power management knobs, such as detecting the amount of energy on-node, the benefits of single cycle power consumption modification, and number of operating modes

These contributions will help enable BSNs achieve it great promise for long term, comprehensive, inexpensive, and unobtrusive monitoring of patients and healthy individuals to run robustly off of harvested energy.

The proposal is organized as follows. Chapter II will investigate the application of PDVS to dynamic workloads. Chapter III will introduce a state of the art battery-less BSN node and explore BSN nodes architectural decisions that enable the BSN node to run off of harvested energy. Chapter IV will propose two revisions of a power manager specific to energy harvesting systems and explore design knobs important to these power managers.

2. Adapting to Dynamic Workloads

2.1 Motivation

Energy efficiency is the most critical metric in modern integrated circuits, setting electronics' functional lifetime. Many systems occasionally require high performance, but their varying workload requirements remain below this upper limit for the majority of their lifetimes. Multimedia processors and medical and environmental sensors are examples of such systems. Designing the system in a static fashion at design time to support this peak performance all the time can substantially increase the total system power, adversely effecting electronics lifetime. Adapting to runtime changes in workload will improve the system's overall energy efficiency. Additionally, by allowing run time changes at a fine grained level (i.e. the component level) will allow for more savings than making run time changes at a global level.

2.2 Prior Art

Two approaches for designing systems and managing their performance at runtime based on dynamic workloads are dynamic voltage scaling (DVS) [1] and multiple supply voltages (multi-V_{DD}) [3]. Both approaches trade off speed for energy when timing slack is available, thus attempting to minimize energy while still meeting any timing constraints. DVS provides the flexibility to alter V_{DD} at runtime, but conventional techniques assign a single voltage to the entire chip (global DVS or single-V_{DD}), shown in Figure 2a, and make voltage changes using an external or internal DC-DC converter, which has transition delays on the order of tens to hundreds of microseconds. Thus, the entire chip must operate at a DVS voltage setting determined by the component with the highest performance requirement, and voltage transitions only can occur when workloads change slowly and infrequently.

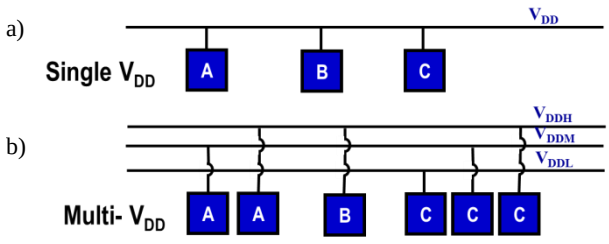

Figure 2. Structure of (a) single V_{DD} and (b) multi-V_{DD}.

Multi-V_{DD}, shown in Figure 2b, statically assigns individual components to voltages to reduce energy for components with timing slack. Operations are then assigned to blocks that are hard tied to a voltage. However, the permanent assignment of V_{DD}s to components limits flexibility to adapt to changing workloads and requirements. If the appropriate voltage component is not available, the operation must be run on a higher voltage component, thus not taking maximum advantage of the timing slack and therefore not maximizing energy efficiency.

Both DVS and multi-V_{DD} do not provide maximum amount of flexibility and maximum amount of granularity to meet peak performance and to achieve maximum energy efficiency when peak performance

is not required. An ideal solution would combine both approaches to meet peak performance and take advantage of timing slack at the block level (i.e. adders, multiplier, etc).

2.3 Panoptic Dynamic Voltage Scaling

To improve energy efficiency, efficiently take advantage of sub-block timing slack, and provide maximum flexibility to adapt to changing workload, [4][5] propose the Panoptic ("all inclusive") Dynamic Voltage Scaling (PDVS) scheme. PDVS provides two main benefits: fine spatial and temporal granularity. Spatial granularity is the ability to assign each component to different voltages at one time. Temporal granularity is the ability to assign the same component to multiple voltages over short time periods.

Figure 3 shows the basic structure of the PDVS architecture. A set of V_{DD} rails are distributed around the chip, and each component, as fine grained as adders and multipliers or as large as cores, can select an operating voltage from one of these rails by connecting to it using PMOS header switches. Header switches improve traditional DVS capabilities by providing rapid and energy efficient transitions between processing voltages on clock edges with minimal energy overhead, rather than relying on DC-DC converters to change their output voltages.

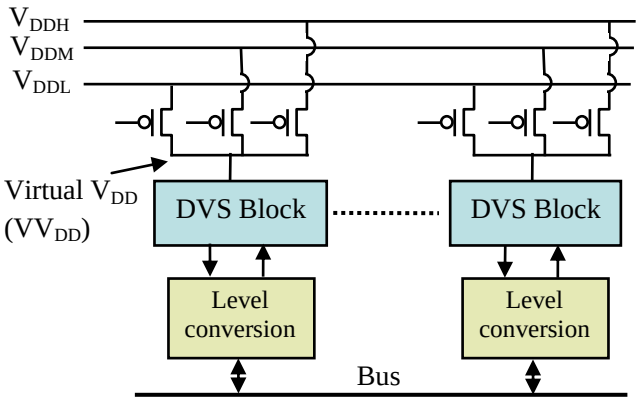

Figure 3. PDVS architecture enabling local fine-grained DVS using header switches and a small set of shared V_{DD}s.

Like multi-V_{DD}, each component can use a voltage that allows it to achieve its required performance without consuming unnecessary additional power independent of other components' voltages. The fine spatial granularity of this architecture gives static savings in energy whenever timing slack exists in the schedule since those components experiencing the slack can switch to lower voltages while still meeting their timing deadlines. PDVS provides this capability with lower area requirements than multi-V_{DD} by assigning voltages to operations rather than components. This means that the same component can execute at different voltages for different operations in the processing flow by switching to the appropriate power rail rather than duplicating the component as is the case with MV_{DD}. The header switches also enhance traditional DVS capabilities by providing rapid and energy efficient transitions between processing voltages on clock edges with minimal energy overhead. PDVS utilizes these header switches instead of relying on external DC-DC converters to adjust V_{DD}, which takes tens to hundreds of µsecs [6]. This allows the capability of rapidly switching from a high performance mode to an ultra low power mode (e.g. using sub-threshold operation) for periods of operation that have extremely relaxed

performance constraints then back to high performance mode again [7]. Prior to this proposed work, no one has explored the system level savings of applying PDVS to a processor. There is a need to understand the system level savings and overheads to motivate its acceptance.

2.4 Research Questions to be Addressed in Dissertation

- Can PDVS be implemented on the system level?
- How does PDVS save energy for three various rate vs. time profiles characteristic to a DSP data flow processor?
- What are the overheads to adding headers and level converter in the PDVS scheme?
- What are the functional limits to V_{DD}-switching speeds?

2.5 Research Questions that will not be Addressed in Dissertation

There are many important questions that are brought up by adapting to runtime variations using PDVS. Since I cannot address them, I recommend the following for future work:

- What is the methodology for selecting voltages for the PDVS scheme?
- Where should the body connection of the headers is tied to?
- What are benefits to using NMOS transistors as headers?
- What are the tradeoffs for sizing headers for PDVS blocks?

2.6 Hypothesis/Thesis

Application of fine grained DVS at runtime using the PDVS scheme will result in higher energy savings on a system level compared to single-V_{DD} and multi-V_{DD} alternatives, despite the overheads of the scheme.

2.7 Approach

We fabricated a 32b data flow processor designed to execute arbitrary data flow graphs (DFGs) at 1 GHz at 1.2V, primarily responsible for digital signal processing. This processor is the first to have PDVS implemented at the system level. Its block diagram is shown in Figure 4. The PDVS data path includes 4 multipliers and 4 adders, each of which uses three header switches to connect to one of three shared V_{DD} rails (V_{DDH}, V_{DDM}, & V_{DDL}). Using more rails provides rapidly diminishing returns. The inputs to the arithmetic components are fed via a programmable crossbar from other computed results, registers, or the 32kb data memory. Level converters are placed at the output of each PDVS component prevent short circuit currents. Each 160b word in the 40kb instruction memory gives control signals (including header gate control) to the data path for one clock step of a DFG. The control structure supports nested looping. This architecture provides flexibility for exploring the benefits of PDVS.

The chip contains three additional data paths for direct comparison with PDVS: single-V_{DD}, multi-V_{DD}, and a PDVS data path optimized to support sub-V_T operation. We will choose lower V_{DD}s that stretch delay by integer numbers of clock cycles. PDVS uses a fixed frequency clock to serve blocks at the highest rate.

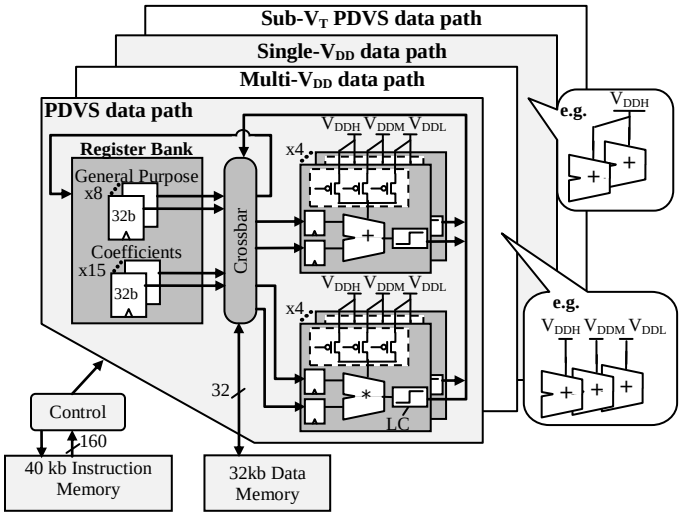

Figure 4. Block diagram of the proposed PDVS data flow processor. SRAMs and control serve four data paths for direct comparison of PDVS with SV_{DD} & MV_{DD}.

This approach allows us to compare PDVS to the other architectures, identify and measure overheads of the PDVS approach, and show the ability to switch into and out of an energy-efficient, sub-threshold on a silicon chip over multiple DFGs.

Through Spectre simulations, we're able to look at the benefits and overheads of different approaches to V_{DD}-switching. The focus of this section is look at short circuit current due to V_{DD}-switching, the energy overheads to switching from a higher voltage to lower voltage, and switching from a lower voltage to a higher voltage as a function of timing of the header gate controls.

Additionally, this work will investigate how PDVS saves energy over various time profiles in Matlab.

2.8 Proposed Contributions

Though PDVS was proposed in [4][5], it has not been applied to applied to large system. This project proposes applying PDVS to a processor to determine PDVS's effectiveness to save energy overall.

The proposed contributions are:

- Demonstrate data flow processor using PDVS in silicon
- Exploration of energy savings compared to single-V_{DD} and multi-V_{DD} alternatives.
- Exploration of PDVS savings for three rate vs. time profiles characteristic to DSP data flow processors.
- Exploration of the overheads and benefits of different speeds and sequences of V_{DD}-switching

2.9 Team and Individual Contribution

The PDVS student team is made up of Saad Arrabi, Kyle Craig, Sudhanshu Khanna, and me. I was responsible for the design and layout of the 32b Baugh Wooley multiplier, the clocking scheme, the control scheme, and investigating multiple design knobs. We all shared in the architectural design of the chip and assisted in simulating and testing the chips. The exploration of PDVS savings for various rate vs time profiles will be done solely by me. More detailed exploration of V_{DD}-switching is also done solely by me.

2.10 Publications

B. H. Calhoun, S. Arrabi, S. Khanna, Y. Shakhsheer, K. Craig, J. Ryan, and J. Lach. "REESES: Rapid Efficient Energy Scalable ElectronicS." *GOMACTech*. March 2010.

S. Khanna, K. Craig, Y. Shakhsheer, S. Arrabi, J. Lach, and B. H. Calhoun. "Stepped Supply Voltage Switching for Energy Constrained Systems." *ISQED*. March 2011.

Y. Shakhsheer, S. Khanna, K. Craig, S. Arrabi, J. Lach, and B. H. Calhoun, "A 90nm Data Flow Processor Demonstrating Fine Grained DVS for Energy Efficient Operation from 0.25V to 1.2V", *CICC*, September 2011.

K. Craig, Y. Shakhsheer, and B. H. Calhoun. "Optimal Power Switch Design for Dynamic Voltage Scaling from High Performance to Subthreshold Operation", *ISLPED* , July 2012.

K. Craig, Y. Shakhsheer, S. Khanna, S. Arrabi, J. Lach, B. H. Calhoun, and S. Kosonocky, "A Programmable Resistive Power Grid for Post-Fabrication Flexibility and Energy Tradeoffs", *ISLPED*, July 2012.

3 BSN Architecture to Run from Harvested Energy

3.1 Motivation

Body sensors networks (BSNs) show great promise for long term, comprehensive, inexpensive, and unobtrusive monitoring of patients and healthy individuals. BSNs address the weaknesses of traditional patient data collection, such as imprecision (qualitative human observation) and under-sampling (infrequent assessment) [8] and help medical personnel diagnose, prevent, and respond to various illnesses such as diabetes, asthma, and heart attacks [9]. Though they show great potential, BSNs have many design challenges that impede their widespread adoption including node lifetime, small form factor for wearability, and cost. Harvesting ambient energy, such as from the sun, vibrations or thermal gradients due to body heat, is an appealing alternative to battery based operation, allowing for an indefinite node lifetime while achieving the required small form factor.

While the amount of harvested energy varies over time, the upper limit of this harvested is in the 10s of µWs. To run solely off this harvested energy, nodes must have an energy efficient architecture to ensure long life time [10]. The node must be able the BSN node must consume low power (10s of µWs) while providing the flexibility to acquire, process and transmit biosignal data while consuming little power.

To be capable of consuming this low amount energy, the node should process data on node to reduce the number of bits transmitted. This allows the power hungry radio to be duty cycled. Appropriate architecture decision such as choosing the correct microcontroller (MCU), bus architecture, and optimizations for the BSN node application will further reduce energy. These architectural design time decisions will allow the node to run off harvested energy and will enable runtime optimizations to adapt to varying amounts of harvested energy.

3.2 Prior Art

Commercial off the shelf (COTS) BSN nodes consume too much power to run solely off of harvested energy. TEMPO3.1, a platform capable of capturing, processing, and wirelessly transmitting six-degrees-of-freedom inertial data, consumes up to 185mW [11]. HealthGear, a real-time wearable system for monitoring, visualizing and analyzing physiological signal, consumes up to 60 mW to sense pulse oximetry [12].

ASIC BSN nodes are capable of achieving lower power than COTS BSN nodes. The authors of [13] present a 0.13-µm CMOS process sub-threshold mixed-signal system-on-chip (SoC) that acquires and processes an ECG signal for wireless ECG monitoring, consuming only 2.6 µW while providing raw ECG data or processed heart rate data. The authors of [14] presents an ECG acquisition and processing SoC with a 3-channel analog front-end (AFE) and flexible, generic DSP components that consumes 12µW. Clock-gating and duty-cycling are used to reduce power, but without voltage reduction techniques (V_{DD}=1.2V), the minimum power is 31.1µW. Both of these SoCs lack radios and a full integrated power delivery system and, therefore, cannot be deployed as standalone systems.

3.3 Research Questions to be Addressed in Dissertation

- Is a BSN node capable of being powered solely by harvested energy?

- What are the benefits of using an on-node custom controller instead of a generic microcontroller?
- How much will optimizing NOPs with a custom controller save at the system level?
- What are the tradeoffs for using circuit switched bus architecture vs. addressable bus architecture for an energy constrained BSN node?

3.4 Research Question that will not be Addressed in Dissertation

There are many important questions that are brought up by making energy efficient BSN architecture decisions. Since I cannot address them, I recommend the following for future work:

- What is the methodology for selecting accelerators for an energy constrained BSN node?
- What is the methodology for selecting the appropriate BSN radio protocol?

3.5 Hypothesis/Thesis

A BSN utilizing a low power microcontroller, low power accelerators, a low power analog front end, and a low power transmitter with intelligent duty cycling will be capable of running solely off harvested energy.

3.6 Approach

We fabricated a 130nm wireless BSN node system-on-chip (SoC) that provides real-time ECG (electrocardiogram), EEG (electroencephalography), and EMG (electromyography) acquisition, signal analysis and processing, and wireless communication over a radio. The node is battery-less, powered entirely from a wearable thermoelectric generator (TEG) with an off-chip storage cap and an on-chip boost converter. This BSN node is the first wireless biosignal processing chip powered solely from a TEG with RF kickstart enabling battery-free operation.

Figure 5 illustrates the highly integrated system block diagram comprising four sections. The node contains the following: an analog front-end with programmable gain and a successive-approximation ADC to amplify and digitize signals as low as a few μVs. An energy harvesting/supply regulation section that boosts an input as low as 30mV up to a regulated 1.35V supply. A subthreshold digital section that performs data processing, mode control, and power management (including power/clock gating of blocks and dynamic voltage scaling (DVS)) based on the available energy on the storage capacitor. The digital section includes a custom digital power management (DPM) processor, general purpose microprocessor (MCU), programmable FIR, 1.5kB instruction SRAM / ROM, 4kB data SRAM, and dedicated accelerators for ECG heart rate (R-R) extraction, atrial fibrillation (AFib) detection, and EMG band energy calculation. Finally, a sub-mW 400/433 MHz MICS/ISM band transmitter that performs BFSK transmission up to 200kbps.

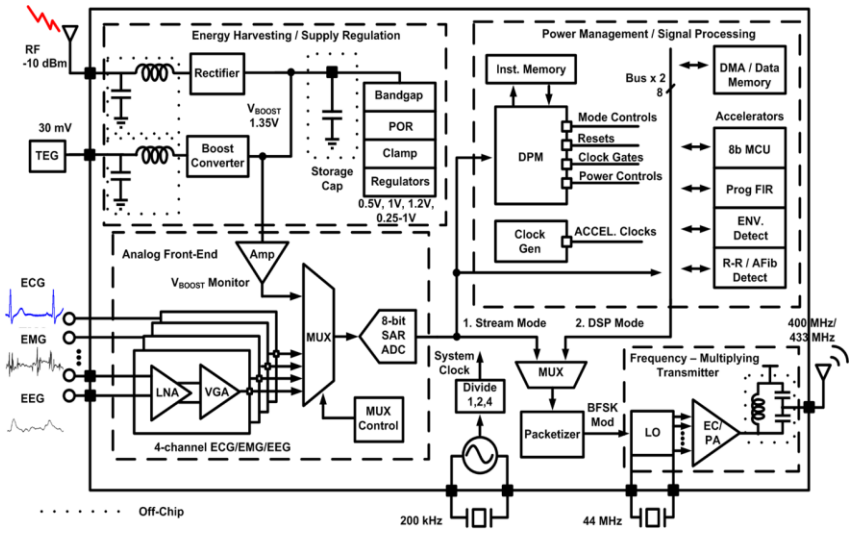

Figure 5. Block diagram of the proposed BSN node revision 1.

The DPM executes arbitrary instructions from an instruction memory and is responsible for managing the AFE, MCU, data memory, accelerator blocks, and transmitter. The DPM provides control signals to the block as show in Figure 3-2. The DPM is capable of duty cycling and power gating power hungry blocks such as individual AFE channels, the transmitter, and its crystal oscillator, resulting in large energy savings.

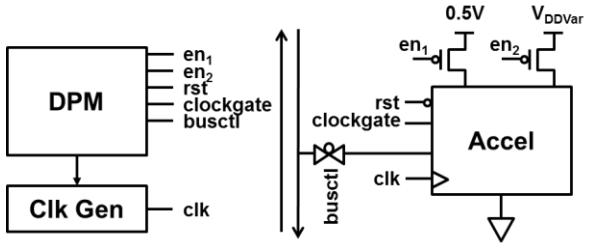

Figure 6. Controls from the DPM to a generic accelerator block.

To save dynamic and leakage energy, the DPM is able to clock gate the whole data memory and individually power gate each bank by asserting bits to the NMOS footer on each bank. Individual accelerators can be clock gated and power gated. Additionally, accelerator blocks are capable of running on a 0.5V supply or a programmable voltage supply. The DPM is capable of setting the voltage of a programmable, variable DC-DC converter (V_{DDVar}) capable of providing 0.3V-1.2V. V_{DDVar} is distributed to each accelerator, allowing blocks to utilize DVS.

The DPM NOPs are optimized to consume ~.25X the energy of a normal DPM instruction. In programs such AFib detection, the DPM spends much of its time idling since the processing is being done in the accelerators for common operating modes. This provides a good opportunity for saving energy. Energy efficient NOPs are run through two DPM instructions, TIMER and STALL. For TIMER, the DPM counts down from one of 16 8b programmed-at-run-time values. STALL is an event based NOP. The node executes NOPs until the next value from the ADC is available, thus not requiring any memory accesses while the DPM awaits the next sample.

This chip will be used to model future BSN chips to help make architecture decisions for future BSN nodes. We will use a custom simulator named SoC Ultra-Low PoweR (SUPR) Model, developed by Alicia Klinefelter, to model nodes. By using data from this chip and this tool, I will be able to explore the benefits and overheads of using a custom controller instead a generic MCU, compare circuit switched bus architecture against an address switched scheme. This exploration will allow designers to make better informed design decisions to build more energy efficient BSN nodes.

3.7 Proposed Contributions

The proposed contributions are:

- A low power, state of the art BSN node capable of running solely off harvested energy
- Exploration of benefits and overheads of using a custom controller (the DPM) in place of a generic microcontroller (PIC-based MCU)
- Comparison of the address scheme vs. circuit switched architecture on the bus level and system level
- Analysis of high power consumption blocks on current chip and future recommendations on how to reduce power for given application

3.8 Team and Individual Contributions

The UVA BSN chip team was Alicia Klinefelter, Jim Boley, Aatmesh Shrivastava, Yanqing Zhang, and me. Our group collaborated with Brian Otis's group from the University of Washington. My individual contribution on the chip was the design and layout of the DPM and with Yanqing, co-designing the chip architecture. My individual contribution to the modeling will be focused on the selection of the microcontroller architecture and bus architecture.

3.9 Publications

F. Zhang, Y. Zhang, J. Silver, Y. Shakhsheer, M. Nagaraju, A. Klinefelter, J. Pandey, J. Boley, E. Carlson, A. Shrivastava, et al., "A Batteryless 19uW MICS/ISM-Band Energy Harvesting Body Area Sensor Node SoC", *ISSCC*, 02/2012.

Y. Zhang, Y. Shakhsheer, A. T. Barth, H. P. C. Jr., S. A. Ridenour, M. A. Hanson, J. Lach, and B. H. Calhoun, "Energy Efficient Design for Body Sensor Nodes", *Journal of Low Power Electronics and Applications*, April 2011.

B. H. Calhoun, Y. Zhang, S. Khanna, K. Craig, Y. Shakhsheer, J. Lach. "A Sub-Threshold FPGA: Energy-Efficient Reconfigurable Logic."*GOMACTech*. March 2011.

4 Adapting to Dynamic Energy Harvesting Rates

4.1 Motivation

Body sensors networks (BSNs) show great promise for long term, comprehensive, inexpensive, and unobtrusive monitoring of patients and healthy individuals. However, they require careful power management due to BSN node's size, weight, lifetime, and cost requirements. The desire for low cost and a small, wearable form factor limits the amount of local energy storage, which is either a battery or a super capacitor. Harvesting ambient energy, such as from the sun, vibrations or thermal gradients due to body heat, is an appealing alternative to battery based operation, allowing for an indefinite node lifetime while achieving the required small form factor. However, energy harvesters' power output is highly environment-dependent. A BSN node capable of running solely of harvested energy require the ability to vary power consumption based the varying amount of harvested energy at runtime. This will extend the life time of the BSN node.

4.2 Prior Art

The authors of [15] recently implemented BSN node that has shown the ability to utilize harvested energy. To sustain operation solely from harvesting, the circuits in the node must consume less power than the power harvesting mechanism produces or temporarily store energy for future use by higher energy operations [16]. There are several issues with using harvested energy. High peak current levels cause significant voltage drops in high-impedance power harvesting sources. Additionally, high power operations, such as transmitting data wirelessly, can consume 100s of μWs [17] and may exceed the power budget set by power harvesting. Lastly, energy harvesters' power output is highly environment-dependent. For example, solar harvesting provides reduced power in cloudy conditions. A power management system requires a power harvesting-specific power management scheme which must account for all these issues to prevent node death (node voltage supply falls below the voltage at which the analog or digital circuits can function). It must adjust node power consumption to the changing power harvesting input to ensure functionality.

Current BSN nodes use a generic microcontroller (MCU) [15] for a variety of functions. The MCU is typically responsible for all on-node processing, power management, and memory accesses. Additionally, MCUs use interrupts to implement data sampling and execute any kind of power management scheme.

Use of a generic MCU to control a power harvesting BSN node is inefficient. Entrusting a large amount of responsibilities to the MCU, such as power management, node control, and data flow, requires overclocking the processor or running at a higher voltage and, thus, consumes more energy. The MCU requires multiple instructions to throttle down the power consumed on the node, thus consuming more energy when energy is limited.

4.3 Research Questions to be Addressed in Dissertation

- What are the benefits to single cycle power modifications?
- What is the methodology for deciding the number of operating modes?
- What are the benefits of adding hysteresis to power management?
- What are the benefits of a stoplight approach?

4.4 Hypothesis/Thesis

A closed loop, energy harvesting-specific stoplight power management, capable of single-cycle power consumption adjustment, will result in longer node lifetime and more robust energy harvesting node. Adding simple hardware to predict the harvesting rate and asymmetric operating mode thresholds will result in a longer node lifetime despite its overheads.

4.5 Approach

Two revisions of an energy harvesting-specific power management systems will be designed and fabrication in a commercial 130nm technology for two BSN nodes. Revision 1 is the first implemented energy harvesting-specific power management system. It features a closed loop power manager capable of changing node power consumption based on the current amount of energy on the capacitor. Revision 2 will build on some of the weaknesses of the previous revision, such as lack of flexibility, and explore some knobs, such as the method for sampling the amount of energy on the storage capacitor and prediction, to build a more robust power manager.

4.5.1 Revision 1

The Digital Power Manager (DPM) is responsible for power management on the battery-less BSN node described in Chapter 3. The DPM's goal is to adjust node power consumption to varying amounts of harvested energy and to extend the BSN node's life. The DPM checks the energy on the storage capacitor and select and implement an operating mode, which limits the maximum amount of node power consumption, in a stoplight fashion. If conditions are poor, the DPM will throttle down processing and/or turn off transmission to save power and preserve the node; likewise, if conditions are favorable, the DPM can allow more blocks to be utilized. Turning off blocks will occur within one clock cycle to conserve the maximum amount of energy.

4.5.1.1 Capacitor Energy

The DPM utilizes the voltage on the storage capacitor to check the energy and make appropriate power management decisions. This voltage corresponds to the stored energy changes with varying node power consumption and power harvesting rate. The capacitor voltage is digitized by the 8b ADC, which also digitizes biosignal. The 8b value is fed directly into the DPM immediately when a new energy value is available.

The capacitor value is only be digitized when a specific DPM instruction is issued. This provides the programmer with the flexibility to bypass any power management restriction if transmitting or processing the data is more important than keeping the node alive and thus, needs to bypass the DPM policies (i.e. the node detected a cardiac arrhythmia and needs to transmit the ECG to notify the doctor over the wireless radio, regardless of energy status).

4.5.1.2 Stoplight

The DPM has three hard-coded operating modes (red, yellow, green). An operating mode is a mode in which only a hardcoded subset of blocks (Table 1) is allowed to be turned on to process or transmit data.

The effectively caps the maximum power consumption for the operating mode. Note that operating modes do not turn blocks on. A DPM must issue an instruction to turn on blocks that are permitted for a given mode. The DPM stoplight issues bits that can prevent mode-restricted blocks from running instructions (Figure 7).

Table 1: Operating Modes

	IMEM	AFE	DMEM	Accel	Transmit
Red	On	Off	Off	Off	Off
Yellow	On	On	On	On	Duty Cycle
Green	On	On	On	On	On

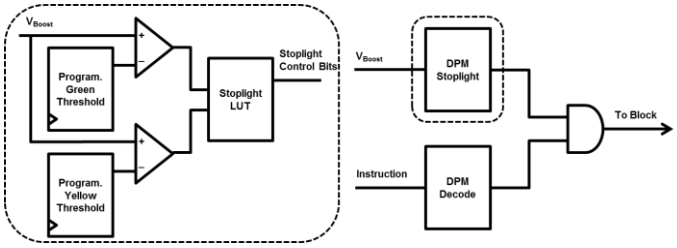

Figure 7. Override structure of the DPM stoplight. The stoplight compares the V_{Boost} value to the threshold, selects the operating mode, and outputs control bits to the chip.

This allows the DPM to modify the node's power consumption within one clock cycle. The DPM's stoplight has the ability to power gate or clock gate individual node blocks, such as the transmitter, accelerator blocks, and analog front end channels to adjust power consumption to adjust to varying amounts of harvested energy. Additional restrictions within the green and yellow modes can be put in place through running a subroutine in code.

4.5.1.3 Threshold Values

Threshold values, programmed through the scan chains at runtime, are used to determine the operating mode. The DPM compares the digitized capacitor value to two 8b threshold values (green and yellow threshold) and sets the operating mode. The current operating mode updates immediately when the stored voltage crosses a threshold. The DPM is capable of jumping from the current operating mode to any mode (i.e. green to red mode) or staying in its current mode at a threshold change.

4.5.2 Revision 2

The goal of this revision is to provide a more flexible power management system capable of being used on multiple energy harvesting BSN nodes. Revision 2 will build on the weaknesses of the previous revision. These weaknesses include lack of flexibility within the programming of operating modes, the inability to scale the number of outputs based on the node, and not allowing the capacitor to fully recover due to the oscillating between operating modes. Revision 2 will explore power management knobs to build a more robust power manager. These knobs include exploring the benefits of using a temperature-invariant ring oscillator to sample the voltage on the capacitor instead of using an ADC, exploring the benefits of single cycle node power reduction, exploring the number of operating modes required for the node, and exploring the benefits and overheads of prediction. These knobs will be explored through Spectre and SUPR model. This revision will be fabricated in a 130nm technology.

4.6 Proposed Contributions

The proposed contributions are:

- The first implemented energy harvesting-specific power management system
- A flexible, energy efficient power manager capable of being used on multiple energy harvesting nodes.
- Exploration of the benefits of single cycle power modification
- Comparison of using a ring oscillator vs. ADC for checking the amount of energy on the storage capacitor
- Exploration of the design decisions for decided the number of operating modes

4.7 Team and Individual Contributions

The UVA BSN chip team was Alicia Klinefelter, Jim Boley, Aatmesh Shrivastava, Yanqing Zhang, and me. Our group collaborated with Brian Otis's group from the University of Washington on revision 1. My individual contributions on both BSN nodes were the design and layout of the power manager and the power management scheme.

4.8 Publications

F. Zhang, Y. Zhang, J. Silver, Y. Shakhsheer, M. Nagaraju, A. Klinefelter, J. Pandey, J. Boley, E. Carlson, A. Shrivastava, et al., "A Batteryless 19uW MICS/ISM-Band Energy Harvesting Body Area Sensor Node SoC", *ISSCC*, 02/2012.

5. Research Tasks

Table 2 below lists the tasks, status, and relevant publications for each research goal.

Table 2: Research task and status

Subject	#	Task Description	Status/Target	Publications
PDVS	1	Design Exploration	Done	
	2	Simulations	Done	
	3	Schematic/Layout	Done	[YAS1]
	4	Test Chip	Aug 2012	[YAS3] [YAS4] [YAS7] [YAS9] [YAS10] [YAS13]
	5	Explore Rate vs. Time Profiles	March 2013	
	6	Explore VDD Switching	March 2013	
BSN Architecture	1	Design Exploration	Done	[YAS5]
	2	RTL	Done	
	3	Synthesis	Done	
	4	Test Chip	Done	[YAS8] [YAS11] [YAS12]
	5	Design Decision Exploration	April 2013	[YAS14]
Power Management	1	Rev 1: Design Exploration	Done	
	2	Rev 1: RTL	Done	
	3	Rev 1: Synthesis	Done	
	4	Rev 1: Test Chip	Done	[YAS11]
	5	Rev 2: Design Exploration	Sept 2012	
	6	Rev 2: RTL	Oct 2012	
	7	Rev 2: Synthesis	Nov 2012	
	8	Rev 2: Test chip	Sept 2013	[YAS15]
Write Up	1	Thesis Writing	Nov 2013	

6. Publications

6.1 Current Publications

[YAS1] B. H. Calhoun, S. Arrabi, S. Khanna, Y. Shakhsheer, K. Craig, J. Ryan, and J. Lach, “REESES: Rapid Efficient Energy Scalable ElectronicS.” *GOMACTech*. March 2010.

[YAS2] J .Stocking, W. Eberhardt, Y. Shakhsheer, J. Paulus, M. Appleby, and B. H. Calhoun, "A Capacitance-Based Whisker-like Artificial Sensor for Fluid Motion Sensing." *IEEE Sensor*. October 2010.

[YAS3] S. Khanna, K. Craig, Y. Shakhsheer, S. Arrabi, J. Lach, and B. H. Calhoun, “Stepped Supply Voltage Switching for Energy Constrained Systems.” *ISQED*. March 2011.

[YAS4] B. H. Calhoun, Y. Zhang, S. Khanna, K. Craig, Y. Shakhsheer, J. Lach, “A Sub-Threshold FPGA: Energy-Efficient Reconfigurable Logic.”*GOMACTech*. March 2011.

[YAS5] Y. Zhang., Y. Shakhsheer, A. T. Barth, H. P. C. Jr., S. A. Ridenour, M. A. Hanson, J. Lach, and B. H. Calhoun, "Energy Efficient Design for Body Sensor Nodes", *Journal of Low Power Electronics and Applications*, April 2011.

[YAS6] W. Eberhardt, Y. Shakhsheer, and B. Calhoun, "A Bio-Inspired Artificial Whisker for Fluid Motion Sensing with Increased Sensitivity and Reliability", *IEEE Sensors*, October 2011.

[YAS7] Y. Shakhsheer, S. Khanna, K. Craig, S. Arrabi, J. Lach, and B. H. Calhoun, "A 90nm Data Flow Processor Demonstrating Fine Grained DVS for Energy Efficient Operation from 0.25V to 1.2V", *CICC*, September 2011.

[YAS8] F. Zhang, Y. Zhang, J. Silver, Y. Shakhsheer, M. Nagaraju, A. Klinefelter, J. Pandey, J. Boley, E. Carlson, A. Shrivastava, B. Otis, and B. H. Calhoun, “A Battery-less 19μW MICS/ISM-Band Energy Harvesting Body Area Sensor Node SoC,” *ISSCC*, February 2012.

[YAS9] K. Craig, Y. Shakhsheer, and B. H. Calhoun. "Optimal Power Switch Design for Dynamic Voltage Scaling from High Performance to Subthreshold Operation," *ISLPED*, July 2012.

[YAS10] K. Craig, Y. Shakhsheer, S. Khanna, S. Arrabi, J. Lach, B. H. Calhoun, and S. Kosonocky, "A Programmable Resistive Power Grid for Post-Fabrication Flexibility and Energy Tradeoffs", *ISLPED*, July 2012.

[YAS11] Y. Shakhsheer, Y. Zhang, B. Otis, and B. H. Calhoun, "A Custom Processor for Node and Power Management of a Battery-less Body Sensor Node in 130nm CMOS", *CICC*, September 2012.

6.2 Anticipated Publications

[YAS11] DPM paper, submitted to CICC

[YAS12] BSN journal paper, invited and submitted to JSSC

[YAS13] PDVS JSSC paper

[YAS14] BSN paper on Revision 2.

[YAS15] Power management Revision 2.

References

[1] E. Carlson, K. Strunz, and B. Otis "A 20 mV Input Boost Converter With Efficient Digital Control for Thermoelectric Energy Harvesting," *JSSC*, Vol. 45, No. 4, April 2010.

[2] A. Chandrakasan, V. Gutnik, and T. Xanthopoulos, "Data driven signal processing: an approach for energy efficient computing," *ISLPED*, Aug. 1996.

[3] J. Howard et al., "A 48-Core IA-32 Message-Passing Processor with DVFS in 45nm CMOS," *ISSCC* , pp. 108-109, Feb. 2010.

[4] M. Putic, L. Di, B. H. Calhoun, and J. Lach, "Panoptic DVS: A Fine-Grained Dynamic Voltage Scaling Framework for Energy Scalable CMOS Design", *ICCD*, pp. 491-497, 01/10/2009.

[5] L. Di, M. Putic, J. Lach, and B. H. Calhoun, "Power Switch Characterization for Fine-Grained Dynamic Voltage Scaling", *ICCD*, pp. 605-611, Aug. 2008.

[6] C. Zheng and D. Ma, "A 10MHz 92.1%-Efficiency Green-Mode Automatic Reconfigurable Switching Converter with Adaptively Compensated Single-Bound Hysteresis Control," *ISSCC*, pp. 204-205, Feb. 2010.

[7] B. Calhoun and A. Chandrakasan. "Ultra-dynamic voltage scaling (UDVS) using sub-threshold operation and local voltage dithering," *JSSC*, Vol. 41, No. 1, pp. 238–245, 2006.

[8] Y. Zhang, Y. Shakhsheer, A. T. Barth, H. P. C. Jr., S. A. Ridenour, M. A. Hanson, J. Lach, and B. H. Calhoun, "Energy Efficient Design for Body Sensor Nodes", *Journal of Low Power Electronics and Applications*, April 2011.

[9] G. Z. Yang, "Body Sensor Networks," London, UK, *Springer-Verlag*, 2006.

[10] G. Chen, S. Hanson, D. Blaauw, D. Sylvester, "Circuit Design Advances for Wireless Sensing Applications," *Proceedings of the IEEE: Special Issue on Wireless Sensor Networks*, Vol. 98, No. 11, pg. 1808 – 1827, Nov. 2010.

[11] A.T. Barth, M.A. Hanson, H.C. Powell, and J. Lach, "TEMPO 3.1: A Body Area Sensor Network Platform for Continuous Movement Assessment," *Sixth International Workshop on Wearable and Implantable Body Sensor Networks, 2009. BSN 2009.* pp.71-76, Jun 2009.

[12] N. Oliver and F, Flores-Mangas, "HealthGear: a real-time wearable system for monitoring and analyzing physiological signals," *International Workshop on Wearable and Implantable Body Sensor Networks, 2006. BSN 2006.* Apr. 2006.

[13] S. Jocke, J. Bolus, S.N. Wooters, A.D. Jurik, A.C. Weaver, T.N. Blalock, and B. H. Calhoun, "A 2.6-μW sub-threshold mixed-signal ECG SoC," *Symposium on VLSI Circuits*, 2009, pp. 60-61.

[14] H. Kim, R. F. Yazicioglu, S. Kim, N. Van Helleputte, A. Artes, M. Konijnenburg, J. Huisken, J. Penders, and C. Van Hoof, "A Configurable and Low-Power Mixed Signal SoC for Portable ECG

Monitoring Applications," *2011 Symposium on VLSI Circuits*, pp. 142-143, Jun. 2011.

[15] G. Chen, M. Fojtik, K. Daeyeon, D. Fick, J. Park, M. Seok, M. Chen, Z. Foo, D. Sylvester, D. Blaauw, "Millimeter-scale nearly perpetual sensor system with stacked battery and solar cells," *ISSCC*, pp. 288-289, Feb. 2010.

[16] B.H. Calhoun, S. Khanna, Y. Zhang, J. Ryan, B. Otis, "System Design Principles Combining Sub-threshold Circuits and Architectures with Energy Scavenging Mechanisms", *ISCAS*, pp. 269-272, 2010.

[17] S. Rai, J. Holleman, J. Pandey, F. Zhang, B. Otis, "A 500μW neural tag with 2μVrms AFE and frequency-multiplying MICS/ISM FSK transmitter," *ISSCC*, pp. 212-213, Feb. 2009.